

Efficient Learning in Large-Scale Combinatorial Semi-Bandits

Zheng Wen

Yahoo Labs, Sunnyvale, CA

ZHENGWEN@YAHOO-INC.COM

Branislav Kveton

Adobe Research, San Jose, CA

KVETON@ADOBE.COM

Azin Ashkan

Technicolor Labs, Los Altos, CA

AZIN.ASHKAN@TECHNICOLOR.COM

Abstract

A stochastic combinatorial semi-bandit is an on-line learning problem where at each step a learning agent chooses a subset of ground items subject to combinatorial constraints, and then observes stochastic weights of these items and receives their sum as a payoff. In this paper, we consider efficient learning in large-scale combinatorial semi-bandits with linear generalization, and as a solution, propose two learning algorithms called *Combinatorial Linear Thompson Sampling* (CombLinTS) and *Combinatorial Linear UCB* (CombLinUCB). Both algorithms are computationally efficient as long as the offline version of the combinatorial problem can be solved efficiently. We establish that CombLinTS and CombLinUCB are also provably statistically efficient under reasonable assumptions, by developing regret bounds that are independent of the problem scale (number of items) and sublinear in time. We also evaluate CombLinTS on a variety of problems with thousands of items. Our experiment results demonstrate that CombLinTS is scalable, robust to the choice of algorithm parameters, and significantly outperforms the best of our baselines.

1. Introduction

Combinatorial optimization is a mature field (Papadimitriou & Steiglitz, 1998), which has countless practical applications. One of the most studied problems in combinatorial optimization is maximization of a modular function

subject to combinatorial constraints. Many important problems, such as minimum spanning tree (MST), shortest path, and maximum-weight bipartite matching, can be viewed as instances of this problem.

In practice, the optimized modular function is often unknown and needs to be learned while repeatedly solving the problem. This class of learning problems was recently formulated as a combinatorial bandit/semi-bandit, depending on the feedback model (Audibert et al., 2014). Since then, many combinatorial bandit/semi-bandit algorithms have been proposed: for the stochastic setting (Gai et al., 2012; Chen et al., 2013; Russo & Van Roy, 2014; Kveton et al., 2015b); for the adversarial setting (Cesa-Bianchi & Lugosi, 2012; Audibert et al., 2014; Neu & Bartók, 2013); and for subclasses of combinatorial problems, matroid and polymatroid bandits (Kveton et al., 2014a;b;c), submodular maximization (Wen et al., 2013; Gabillon et al., 2013), and cascading bandits (Kveton et al., 2015a). Many regret bounds have been established for the combinatorial semi-bandit algorithms. To achieve an $O(\sqrt{n})$ dependence on time n , all of the regret bounds are $\Omega(\sqrt{L})$, where L is the number of items. The dependence on L is intrinsic because the algorithms estimate the weight of each item separately, and matching lower bounds have been established (Section 3.2).

However, in many real-world problems, the number of items L is intractably large. For instance, online advertising in a mainstream commercial website can be viewed as a bipartite matching problem with millions of users and products; routing in the Internet can be formulated as a shortest path problem with billions of edges. Thus, learning algorithms with $\Omega(\sqrt{L})$ regret are impractical in such problems. On the other hand, in many problems, items have features and their weights are similar when the features are similar. In movie recommendation, for instance, the expected ratings of movies that are close in the latent space are also similar. In this work, we show how to leverage this struc-

ture to learn to make good decisions more efficiently. More specifically, we assume a *linear generalization* across the items: conditioned on the features of an item, the expected weight of that item can be estimated using a linear model. Our goal is to develop more efficient learning algorithms for combinatorial semi-bandits with linear generalization.

It is relatively easy to extend many linear bandit algorithms, such as Thompson sampling (Thompson, 1933; Agrawal & Goyal, 2012; Russo & Van Roy, 2013) and Linear UCB (LinUCB, see Auer (2002); Dani et al. (2008); Abbasi-Yadkori et al. (2011)), to combinatorial semi-bandits with linear generalization. In this paper, we propose two learning algorithms, Combinatorial Linear Thompson Sampling (CombLinTS) and Combinatorial Linear UCB (CombLinUCB), based on Thompson sampling and LinUCB. Both CombLinTS and CombLinUCB are computationally efficient, as long as the offline version of the combinatorial problem can be solved efficiently. The first major contribution of the paper is that we establish a *Bayes regret bound* on CombLinTS and a *regret bound* on CombLinUCB, under reasonable assumptions. Both bounds are L -independent, and sublinear in time. The second major contribution of the paper is that we evaluate CombLinTS on a variety of problems with thousands of items, and two of these problems are based on real-world datasets. We only evaluate CombLinTS since recent literature (Chapelle & Li, 2011) suggests that Thompson sampling algorithms usually outperform UCB-like algorithms in practice. Our experimental results demonstrate that CombLinTS is scalable, robust to the choice of algorithm parameters, and significantly outperforms the best of our baselines. It is worth mentioning that our derived L -independent regret bounds also hold in cases with $L = \infty$. Moreover, as we will discuss in Section 7, our proposed algorithms and their analyses can be easily extended to the *contextual combinatorial semi-bandits*.

Finally, we briefly review some relevant papers. Gabillon et al. (2014) and Yue & Guestrin (2011) focus on submodular maximization with linear generalization. Our paper differs from these two papers in the following two aspects: (1) our paper allows general combinatorial constraints while they do not; (2) our paper focuses on maximization of modular functions while they focus on submodular maximization.

2. Combinatorial Optimization

We focus on a class of combinatorial optimization problems that aim to find a *maximum-weight* set from a given family of sets. Specifically, one such combinatorial optimization problem can be represented as a triple $(E, \mathcal{A}, \mathbf{w})$, where (1) $E = \{1, \dots, L\}$ is a set of L items, called the *ground set*, (2) $\mathcal{A} \subseteq \{A \subseteq E : |A| \leq K\}$ is a family of

subsets of E with up to K items, where $K \leq L$, and (3) $\mathbf{w} : E \rightarrow \mathbb{R}$ is a *weight function* that assigns each item e in the ground set E a real number. The total weight of all items in a set $A \subseteq E$ is defined as:

$$f(A, \mathbf{w}) = \sum_{e \in A} \mathbf{w}(e), \quad (1)$$

which is a linear functional of \mathbf{w} and a modular function in A . A set A^{opt} is a maximum-weight set in \mathcal{A} if:

$$A^{\text{opt}} \in \arg \max_{A \in \mathcal{A}} f(A, \mathbf{w}) = \arg \max_{A \in \mathcal{A}} \sum_{e \in A} \mathbf{w}(e). \quad (2)$$

Many classical combinatorial optimization problems, such as finding an MST, bipartite matching, the shortest path problem and the traveling salesman problem (TSP), have form (2). Though some of these problems can be solved efficiently (e.g. bipartite matching), others (e.g. TSP) are known to be NP-hard. However, for many such NP-hard problems, there exist computationally efficient *approximation algorithms* and/or *randomized algorithms* that achieve near-optimal solutions with high probability. Similarly to Chen et al. (2013), in this paper, we allow the agent to use any approximation / randomized algorithm ORACLE to solve (2), and denote its solution as $A^* = \text{ORACLE}(E, \mathcal{A}, \mathbf{w})$. To distinguish from a learning algorithm, we refer to a combinatorial optimization algorithm as an *oracle* in this paper.

3. Combinatorial Semi-Bandits with Linear Generalization

Many real-world problems are combinatorial in nature. In recommender systems, for instance, the user is typically recommended K items out of L . The value of an item, such as the expected rating of a movie, is never known perfectly and has to be refined while repeatedly recommending to the pool of the users. Recommender problems are known to be highly structured. In particular, it is well known that the user-item matrix is typically low-rank (Koren et al., 2009) and that the value of an item can be written as a linear combination of its position in the latent space. In this work, we propose a learning algorithm for combinatorial optimization that leverages this structure. In particular, we assume that the weight of each item is a linear function of its features and then we learn the parameters of this model, jointly for all items.

3.1. Combinatorial Semi-Bandits

We formalize our learning problem as a combinatorial semi-bandit. A combinatorial semi-bandit is a triple (E, \mathcal{A}, P) , where E and \mathcal{A} are defined in Section 2 and P is a probability distribution over the weights $\mathbf{w} \in \mathbb{R}^L$ of the items in the ground set E . We assume that the weights \mathbf{w} are drawn i.i.d. from P . The mean weight is denoted by

$\bar{\mathbf{w}} = \mathbb{E}[\mathbf{w}]$. Each item e is associated with an *arm* and we assume that *multiple arms* can be pulled. A subset of arms $A \subseteq E$ can be pulled if and only if $A \in \mathcal{A}$. The return of pulling arms A is $f(A, \mathbf{w})$ (Equation (1)), the sum of the weights of all items in A . After the arms A are pulled, we observe the individual return of each arm, $\{\mathbf{w}(e) : e \in A\}$. This feedback model is known as *semi-bandit* (Audibert et al., 2014).

We assume that the combinatorial structure (E, \mathcal{A}) is known and the distribution P is unknown. We would like to stress that we do not make any structural assumptions on P . The optimal solution to our problem is a maximum-weight set in expectation:

$$A^{\text{opt}} \in \arg \max_{A \in \mathcal{A}} \mathbb{E}_{\mathbf{w}}[f(A, \mathbf{w})] = \arg \max_{A \in \mathcal{A}} \sum_{e \in A} \bar{\mathbf{w}}(e). \quad (3)$$

This objective is equivalent to the one in Equation (2).

Our learning problem is episodic. In each episode t , the learning agent adaptively chooses $A^t \in \mathcal{A}$ based on its observations of the weights up to episode t , gains $f(A^t, \mathbf{w}_t)$, and observes the weights of all chosen items in episode t , $\{(e, \mathbf{w}_t(e)) : e \in A^t\}$. The learning agent interacts with the combinatorial semi-bandit for n times and its goal is to maximize the expected cumulative return in n -episodes $\mathbb{E}[\sum_{t=1}^n f(A^t, \mathbf{w}_t)]$, where the expectation is over (1) the random weights \mathbf{w}_t 's, (2) possible randomization in the learning algorithm, and (3) $\bar{\mathbf{w}}$ if it is randomly generated. Notice that the choice of A^t impacts both the return and observations in episode t . So we need to trade off *exploration* and *exploitation*, similarly to other bandit problems.

3.2. Linear Generalization

As we have discussed in Section 1, many provably efficient algorithms have been developed for various combinatorial semi-bandits of form (3) (Chen et al., 2013; Gai et al., 2012; Russo & Van Roy, 2014; Kveton et al., 2014b; 2015b). However, since there are L parameters to learn and these algorithms do not consider *generalization* across items, the derived upper bounds on the expected cumulative regret and/or the Bayes cumulative regret of these algorithms are at least $O(\sqrt{L})$. Furthermore, Audibert et al. (2014) has derived an $\Omega(\sqrt{LK n})$ lower bound on adversarial combinatorial semi-bandits, while Kveton et al. (2014b; 2015b) have derived asymptotic $\Omega(L \log(n)/\Delta)$ gap-dependent lower bounds on stochastic combinatorial semi-bandits, where Δ is an appropriate ‘‘gap’’.

However, in many modern combinatorial semi-bandit problems, L tends to be enormous. Thus, an $O(\sqrt{L})$ regret is unacceptably large in these problems. On the other hand, in many practical problems, there exists a *generalization model* based on which the weight of one item can be (approximately) inferred based on the weights of

other items. By exploiting such generalization models, an $o(\sqrt{L})$ or even an L -independent cumulative regret might be achieved.

In this paper, we assume that there is a (possibly imperfect) linear generalization model across the items. Specifically, we assume that the agent knows a *generalization matrix* $\Phi \in \mathbb{R}^{L \times d}$ s.t. $\bar{\mathbf{w}}$ either lies in or is ‘‘close’’ to the subspace $\text{span}[\Phi]$. We use ϕ_e to denote the transpose of the e -th row of Φ , and refer to it as the *feature vector* of item e . Without loss of generality, we assume that $\text{rank}[\Phi] = d$.

Similar to some existing literature (Wen & Van Roy, 2013; Van Roy & Wen, 2014), we distinguish between the *coherent learning* cases, in which $\bar{\mathbf{w}} \in \text{span}[\Phi]$, and the *agnostic learning* cases, in which $\bar{\mathbf{w}} \notin \text{span}[\Phi]$. Like existing literature on linear bandits (Dani et al., 2008; Abbasi-Yadkori et al., 2011), the analysis in this paper focuses on coherent learning cases. However, we would like to emphasize that both of our proposed algorithms, CombLinTS and CombLinUCB, are also applicable to the agnostic learning cases. As is demonstrated in Section 6, CombLinTS performs well in the agnostic learning cases.

Finally, we define $\theta^* = \arg \min_{\theta} \|\bar{\mathbf{w}} - \Phi\theta\|_2$. Since $\text{rank}[\Phi] = d$, θ^* is uniquely defined. Moreover, in coherent learning cases, we have $\bar{\mathbf{w}} = \Phi\theta^*$.

3.3. Performance Metrics

Let $A^* = \text{ORACLE}(E, \mathcal{A}, \bar{\mathbf{w}})$. In this paper, we measure the performance loss of a learning algorithm with respect to A^* . Recall that the learning algorithm chooses A^t in episode t , we define $R_t = f(A^*, \mathbf{w}_t) - f(A^t, \mathbf{w}_t)$ as the *realized regret* in episode t . If the expected weight $\bar{\mathbf{w}}$ is fixed but unknown, we define the *expected cumulative regret* of the learning algorithm in n episodes as

$$R(n) = \sum_{t=1}^n \mathbb{E}[R_t | \bar{\mathbf{w}}], \quad (4)$$

where the expectation is over random weights and possible randomization in the learning algorithm. If necessary, we denote $R(n)$ as $R(n; \bar{\mathbf{w}})$ to emphasize the dependence on $\bar{\mathbf{w}}$. On the other hand, if $\bar{\mathbf{w}}$ is randomly generated or the agent has a prior belief in $\bar{\mathbf{w}}$, then from Russo & Van Roy (2013), the *Bayes cumulative regret* of the learning algorithm in n episodes is defined as

$$R_{\text{Bayes}}(n) = \mathbb{E}_{\bar{\mathbf{w}}}[R(n; \bar{\mathbf{w}})] = \sum_{t=1}^n \mathbb{E}[R_t], \quad (5)$$

where the expectation is also over $\bar{\mathbf{w}}$. That is, $R_{\text{Bayes}}(n)$ is a weighted average of $R(n; \bar{\mathbf{w}})$ under the prior on $\bar{\mathbf{w}}$.

4. Learning Algorithms

In this section, we propose two learning algorithms for combinatorial semi-bandits: Combinatorial Linear Thompson Sampling (CombLinTS) and Combinatorial Linear

Algorithm 2 Combinatorial Linear Thompson Sampling

Input: Combinatorial structure (E, \mathcal{A}) , generalization matrix $\Phi \in \mathbb{R}^{L \times d}$, algorithm parameters $\lambda, \sigma > 0$, oracle ORACLE

Initialize $\Sigma_1 \leftarrow \lambda^2 I \in \mathbb{R}^{d \times d}$ and $\bar{\theta}_1 = 0 \in \mathbb{R}^d$

for all $t = 1, 2, \dots, n$ **do**

 Sample $\theta_t \sim N(\bar{\theta}_t, \Sigma_t)$

 Compute $A^t \leftarrow \text{ORACLE}(E, \mathcal{A}, \Phi \theta_t)$

 Choose set A^t , and observe $\mathbf{w}_t(e), \forall e \in A^t$

 Compute $\bar{\theta}_{t+1}$ and Σ_{t+1} based on Algorithm 1

end for

UCB (CombLinUCB), which are respectively motivated by Thompson sampling and LinUCB. Both algorithms maintain a mean vector $\bar{\theta}_t$ and a covariance matrix Σ_t , and use Kalman filtering to update $\bar{\theta}_t$ and Σ_t . They differ in how to choose A^t (i.e. how to explore) in each episode t : CombLinTS chooses A^t based on a randomly sampled coefficient vector θ_t , while CombLinUCB chooses A^t based on the *optimism in the face of uncertainty (OFU)* principle.

4.1. Combinatorial Linear Thompson Sampling

The pseudocode of CombLinTS is given in Algorithm 2, where (E, \mathcal{A}) is the combinatorial structure, Φ is the generalization matrix, ORACLE is a combinatorial optimization algorithm, and λ and σ are two algorithm parameters controlling the *learning rate*. Specifically, λ is an *inverse-regularization* parameter and smaller λ makes the covariance matrix Σ_t closer to 0. Thus, a too small λ will lead to insufficient exploration and significantly reduce the performance of CombLinTS. On the other hand, σ controls the decrease rate of the covariance matrix Σ_t . In particular, a large σ will lead to slow learning, while a too small σ will make the algorithm quickly converge to some sub-optimal coefficient vector.

In each episode t , Algorithm 2 consists of three steps. First, it randomly samples a coefficient vector θ_t from a Gaussian distribution. Second, it computes A^t based on θ_t and the pre-specified oracle. Finally, it updates the mean vector $\bar{\theta}_{t+1}$ and the covariance matrix Σ_{t+1} based on Kalman filtering (Algorithm 1).

It is worth pointing out that if (1) $\bar{\mathbf{w}} = \Phi \theta^*$, (2) the prior on θ^* is $N(0, \lambda^2 I)$, and (3) $\forall (t, e)$, the noise $\eta_t(e) = \mathbf{w}_t(e) - \bar{\mathbf{w}}(e)$ is independently sampled from $N(0, \sigma^2)$, then in each episode t , the CombLinTS algorithm samples θ_t from the posterior distribution of θ^* . We henceforth refer to a case satisfying condition (1)-(3) as a *coherent Gaussian case*. Obviously, the CombLinTS algorithm can be applied to more general cases, even to cases with no prior and/or agnostic learning cases.

Algorithm 3 Combinatorial Linear UCB

Input: Combinatorial structure (E, \mathcal{A}) , generalization matrix $\Phi \in \mathbb{R}^{L \times d}$, algorithm parameters $\lambda, \sigma, c > 0$, oracle ORACLE

Initialize $\Sigma_1 \leftarrow \lambda^2 I \in \mathbb{R}^{d \times d}$ and $\bar{\theta}_1 = 0 \in \mathbb{R}^d$

for all $t = 1, 2, \dots, n$ **do**

 Define the UCB weight vector $\hat{\mathbf{w}}_t$ as

$$\hat{\mathbf{w}}_t(e) = \langle \phi_e, \bar{\theta}_t \rangle + c \sqrt{\phi_e^T \Sigma_t \phi_e} \quad \forall e \in E$$

 Compute $A^t \leftarrow \text{ORACLE}(E, \mathcal{A}, \hat{\mathbf{w}}_t)$

 Choose set A^t , and observe $\mathbf{w}_t(e), \forall e \in A^t$

 Compute $\bar{\theta}_{t+1}$ and Σ_{t+1} based on Algorithm 1

end for

4.2. Combinatorial Linear UCB

The pseudocode of CombLinUCB is given in Algorithm 3, where E, \mathcal{A}, Φ and ORACLE are defined the same as in Algorithm 2, and λ, σ , and c are three algorithm parameters. Similarly, λ is an inverse-regularization parameter, σ controls the decrease rate of the covariance matrix, and c controls the *degree of optimism* (exploration). Specifically, if c is too small, the algorithm might converge to some sub-optimal coefficient vector due to insufficient exploration; on the other hand, too large c will lead to excessive exploration and slow learning.

In each episode t , Algorithm 3 also consists of three steps. First, for each $e \in E$, it computes an upper confidence bound (UCB) $\hat{\mathbf{w}}_t(e)$. Second, it computes A^t based on $\hat{\mathbf{w}}_t$ and the pre-specified oracle. Finally, it updates $\bar{\theta}_{t+1}$ and Σ_{t+1} based on Kalman filtering (Algorithm 1).

5. Regret Bounds

In this section, we present a Bayes regret bound on CombLinTS, and a regret bound on CombLinUCB. We will also briefly discuss how these bounds are derived, as well as their tightness. The detailed proofs are left to the appendices. Without loss of generality, throughout this section, we assume that $\|\phi_e\|_2 \leq 1, \forall e \in E$.

5.1. Bayes Regret Bound on CombLinTS

We have the following upper bound on $R_{\text{Bayes}}(n)$ when CombLinTS is applied to a coherent Gaussian case with the right parameter.

Theorem 1. *If (1) $\bar{\mathbf{w}} = \Phi \theta^*$, (2) the prior on θ^* is $N(0, \lambda^2 I)$, (3) the noises are i.i.d. sampled from $N(0, \sigma^2)$, and (4) $\lambda \geq \sigma$, then under CombLinTS algorithm with pa-*

Algorithm 1 Compute $\bar{\theta}_{t+1}$ and Σ_{t+1} based on Kalman Filtering

Input: $\bar{\theta}_t, \Sigma_t, \sigma$, and feature-observation pairs $\{(\phi_e, \mathbf{w}_t(e)) : e \in A^t\}$

 Initialize $\bar{\theta}_{t+1} \leftarrow \bar{\theta}_t$ and $\Sigma_{t+1} \leftarrow \Sigma_t$
for $k = 1, \dots, |A^t|$ **do**

 Update $\bar{\theta}_{t+1}$ and Σ_{t+1} as follows, where a_k^t is the k th element in A^t

$$\bar{\theta}_{t+1} \leftarrow \left[I - \frac{\Sigma_{t+1} \phi_{a_k^t} \phi_{a_k^t}^T}{\phi_{a_k^t}^T \Sigma_{t+1} \phi_{a_k^t} + \sigma^2} \right] \bar{\theta}_{t+1} + \left[\frac{\Sigma_{t+1} \phi_{a_k^t}}{\phi_{a_k^t}^T \Sigma_{t+1} \phi_{a_k^t} + \sigma^2} \right] \mathbf{w}_t(a_k^t) \quad \text{and} \quad \Sigma_{t+1} \leftarrow \Sigma_{t+1} - \frac{\Sigma_{t+1} \phi_{a_k^t} \phi_{a_k^t}^T \Sigma_{t+1}}{\phi_{a_k^t}^T \Sigma_{t+1} \phi_{a_k^t} + \sigma^2},$$

end for
Output: $\bar{\theta}_{t+1}$ and Σ_{t+1}

parameter (Φ, λ, σ) , we have

$$R_{\text{Bayes}}(n) \leq \tilde{O} \left(K \lambda \sqrt{dn \min \{\ln(L), d\}} \right). \quad (6)$$

Notice that condition (1)-(3) ensure it is a coherent Gaussian case, and condition (4) almost always holds¹. The \tilde{O} notation hides the logarithm factors. We also note that Equation (6) is a minimum of two bounds. The first bound is L -dependent, but it is only $O(\sqrt{\ln(L)})$; on the other hand, the second bound is L -independent, but is $\tilde{O}(d)$ instead of $\tilde{O}(\sqrt{d})$. We would like to emphasize that Theorem 1 holds even if ORACLE is an approximation/randomized algorithm.

We now outline the proof of Theorem 1, which is motivated by Russo & Van Roy (2013) and Dani et al. (2008). Let \mathcal{H}_t denote the ‘‘history’’ (i.e. all the available information) by the start of episode t . Note that from the Bayesian perspective, conditioning on \mathcal{H}_t , θ^* and θ_t are i.i.d. drawn from $N(\bar{\theta}_t, \Sigma_t)$ (Russo & Van Roy, 2013). This is because that conditioning on \mathcal{H}_t , the posterior belief in θ^* is $N(\bar{\theta}_t, \Sigma_t)$ and based on Algorithm 2, θ_t is independently sampled from $N(\bar{\theta}_t, \Sigma_t)$. Since ORACLE is a fixed combinatorial optimization algorithm (even though it can be independently randomized), and E, \mathcal{A}, Φ are all fixed, then conditioning on \mathcal{H}_t , A^* and A^t are also i.i.d., furthermore, A^* is conditionally independent of θ_t , and A^t is conditionally independent of θ^* .

To simplify the exposition, $\forall \theta \in \mathbb{R}^d$ and $\forall A \subseteq E$, we define

$$g(A, \theta) = \sum_{e \in A} \langle \phi_e, \theta \rangle,$$

where $\langle \cdot, \cdot \rangle$ is an alternative notation for inner product. Thus we have $\mathbb{E}[R_t | \mathcal{H}_t] = \mathbb{E}[g(A^*, \theta^*) - g(A^t, \theta^*) | \mathcal{H}_t]$. We also define a UCB function $U_t : 2^E \rightarrow \mathbb{R}$ as

$$U_t(A) = \sum_{e \in A} \left[\langle \phi_e, \bar{\theta}_t \rangle + c \sqrt{\phi_e^T \Sigma_t \phi_e} \right],$$

¹Condition (4) is not essential, please refer to Theorem 3 in Appendix A for a Bayes regret bound without condition (4).

where $c > 0$ is a constant to be specified. Notice that conditioning on \mathcal{H}_t , U_t is a deterministic function and A^*, A^t are i.i.d., then $\mathbb{E}[U_t(A^t) - U_t(A^*) | \mathcal{H}_t] = 0$ and

$$\begin{aligned} \mathbb{E}[R_t | \mathcal{H}_t] &= \mathbb{E}[g(A^*, \theta^*) - U_t(A^*) | \mathcal{H}_t] \\ &\quad + \mathbb{E}[U_t(A^t) - g(A^t, \theta^*) | \mathcal{H}_t]. \end{aligned} \quad (7)$$

Theorem 1 follows by respectively bounding the two terms on the righthand side of Equation (7). Two key observations are (1) if $c = \tilde{O} \left(\sqrt{\min \{\ln(L), d\}} \right)$, then

$$\mathbb{E}[g(A^*, \theta^*) - U_t(A^*) | \mathcal{H}_t] = O(1),$$

and (2)

$$\mathbb{E}[U_t(A^t) - g(A^t, \theta^*) | \mathcal{H}_t] = c \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \middle| \mathcal{H}_t \right],$$

and we have a worst-case bound (see Lemma 4 in Appendix A) on $\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e}$. Please refer to Appendix A for the detailed proof for Theorem 1.

Finally, we briefly discuss the tightness of our bound. Without loss of generality, we assume that $\lambda = 1$. For the special case when $\Phi = I$ (i.e. no generalization), Russo & Van Roy (2014) provides an $O(\sqrt{LK \log(L/K)n})$ upper bound on $R_{\text{Bayes}}(n)$ when Thompson sampling is applied, and Audibert et al. (2014) provides an $\Omega(\sqrt{LK n})$ lower bound². Since $L = d$ when $\Phi = I$, the above results indicate that for general Φ , the best upper bound one can hope is $O(\sqrt{K d n})$. Hence, our bound is at most $\tilde{O}(\sqrt{K \min \{\ln(L), d\}})$ larger. It is well-known that the $O(\sqrt{d})$ factor is due to linear generalization (Dani et al., 2008; Abbasi-Yadkori et al., 2011), and as is discussed in the appendix (see Remark 1), the extra $O(\sqrt{K})$ factor is also due to linear generalization. They might be intrinsic, but we leave the final word and tightness analysis to future work.

²Audibert et al. (2014) focuses on the adversarial setting but the lower bound is stochastic. So it is a reasonable lower bound to compare with.

5.2. Regret Bound on CombLinUCB

Under the assumptions that (1) the support of P is a subset of $[0, 1]^L$ (i.e. $\mathbf{w}_t(e) \in [0, 1] \forall t$ and $\forall e \in E$), and (2) the oracle ORACLE *exactly* solves the offline optimization problem³, we have the following upper bound on $R(n)$ when CombLinUCB is applied to coherent learning cases:

Theorem 2. *For any $\lambda, \sigma > 0$, any $\delta \in (0, 1)$, and any c satisfying*

$$c \geq \frac{1}{\sigma} \sqrt{d \ln \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right)} + 2 \ln \left(\frac{1}{\delta} \right) + \frac{\|\theta^*\|_2}{\lambda}, \quad (8)$$

if $\bar{\mathbf{w}} = \Phi\theta^$ and the above two assumptions hold, then under CombLinUCB algorithm with parameter $(\Phi, \lambda, \sigma, c)$, we have*

$$R(n) \leq 2cK\lambda \sqrt{\frac{dn \ln \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right)}{\ln \left(1 + \frac{\lambda^2}{\sigma^2} \right)}} + nK\delta.$$

Generally speaking, the proof for Theorem 2 proceeds as follows. We first construct a confidence set G of θ^* based on the “self normalized bound” developed in Abbasi-Yadkori et al. (2011). Then we decompose the regret over the high-probability “good” event G and the low-probability “bad” event \bar{G} , where \bar{G} is the complement of G . Finally, we bound the term associated with the event G based on the same worst-case bound on $\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e}$ used in the analysis for CombLinTS (see Lemma 4 in Appendix A), and bound the term associated with the event \bar{G} based on a naive bound. Please refer to Appendix B for the detailed proof of Theorem 2.

Notice that if we choose $\lambda = \sigma = 1$, $\delta = 1/(nK)$, and c as the lower bound specified in Inequality (8), then the regret bound derived in Theorem 2 is also $\tilde{O}(Kd\sqrt{n})$. Compared with the lower bound derived in Audibert et al. (2014), this bound is at most $\tilde{O}(\sqrt{Kd})$ larger. Similarly, the extra $O(\sqrt{K})$ and $O(\sqrt{d})$ factors are also due to linear generalization.

Finally, we would like to clarify that the assumption that the support of P is bounded is not essential. By slightly modifying the analysis, we can achieve a similar high-probability bound on the *realized cumulative regret* as long as P is *sub-Gaussian*. We also want to point out that the L -independent bounds derived in both Theorem 1 and 2 will still hold even if $L = \infty$.

6. Experiments

In this section, we evaluate CombLinTS on three problems. The first problem is synthetic, but the last two problems

are constructed based on real-world datasets. As we have discussed in Section 1, we only evaluate CombLinTS since in practice Thompson sampling algorithms usually outperform the UCB-like algorithms. Our experiment results in the synthetic problem demonstrate that CombLinTS is both scalable and robust to the choice of algorithm parameters. They also suggest the Bayes regret bound derived in Theorem 1 is likely to be tight. On the other hand, our experiment results in the last two problems show the *value of linear generalization* in real-world settings: with domain-specific but imperfect linear generalization (i.e. agnostic learning), CombLinTS can significantly outperform state-of-the-art learning algorithms that do not exploit linear generalization, which serve as baselines in these two problems.

In all three problems, the oracle ORACLE exactly solves the offline combinatorial optimization problem. Moreover, in the two real-world problems, we demonstrate the experiment results using a new performance metric, the *expected per-step return* in n episodes, which is defined as

$$\frac{1}{n} \mathbb{E}_{\mathbf{w}_1, \dots, \mathbf{w}_n} [\sum_{t=1}^n f(A^t, \mathbf{w}_t) | \bar{\mathbf{w}}]. \quad (9)$$

Obviously, it is the expected cumulative return in n episodes divided by n . We demonstrate experiment results using expected cumulative return rather than $R(n)$ since it is more illustrative.

6.1. Longest Path

We first evaluate CombLinTS on a synthetic problem. Specifically, we experiment with a stochastic longest path problem on an $(m+1) \times (m+1)$ square grid⁴. The items in the ground set E are the edges in the grid, $L = 2m(m+1)$ in total. The feasible set \mathcal{A} are all paths in the grid from the upper left corner to the bottom right corner that follow the directions of the edges. The length of these paths is $K = 2m$. In this problem, we focus on coherent Gaussian cases and randomly sample the linear generalization matrix $\Phi \in \mathbb{R}^{L \times d}$ to weaken the dependence on a particular choice of Φ .

Our experiments are parameterized by a sextuple $(m, d, \lambda_{\text{true}}, \sigma_{\text{true}}, \lambda, \sigma)$, where m , d , λ , and σ are defined before and λ_{true} and σ_{true} are respectively the true standard deviations of θ^* and the observation noises. In each round of simulation, we first construct a problem instance as follows: (1) generate Φ by sampling each component of Φ i.i.d. from $N(0, 1)$; (2) sample θ^* independently from $N(0, \lambda_{\text{true}}^2 I)$ and set $\bar{\mathbf{w}} = \Phi\theta^*$; and (3) $\forall(t, e)$, the observation noise $\eta_t(e) = \mathbf{w}_t(e) - \bar{\mathbf{w}}(e)$ is i.i.d. sampled from $N(0, \sigma_{\text{true}}^2)$. Then we apply CombLinTS

³If ORACLE is an approximation algorithm, a variant of Theorem 2 can be proved (see Appendix D).

⁴That is, each side has m edges and $m+1$ nodes. Notice that the longest path problem and the shortest path problem are mathematically equivalent.




Figure 1. Experiment results in the longest path problem

with parameter (λ, σ) to the constructed instance for n episodes. Notice that in general $(\lambda, \sigma) \neq (\lambda_{\text{true}}, \sigma_{\text{true}})$. We average the experiment results over 200 simulations to estimate the Bayes cumulative regret $R_{\text{Bayes}}(n)$.

We start with a “default case” with $m = 30$, $d = 200$, $\lambda_{\text{true}} = \lambda = 10$ and $\sigma_{\text{true}} = \sigma = 1$. Notice in this case $L = 1860$ and $|\mathcal{A}| \approx 1.18 \times 10^{17}$. We choose $n = 150$ since in the default case, the Bayes per-episode regret of CombLinTS vanishes far before period 150. In the default case $R_{\text{Bayes}}(150) \approx 1.56 \times 10^4$. In the experiments, we vary only one and only one parameter while keeping all the other parameters fixed to their “default values” specified above to demonstrate the *scalability* and *robustness* of CombLinTS.

First, we study how the Bayes cumulative regret of CombLinTS scales with the size of the problem by varying $m = 10, 20, \dots, 100$, and show the result in Figure 1(a). The experiment results show that $R_{\text{Bayes}}(150)$ roughly increases linearly with m , which indicates that CombLinTS is scalable with respect to the problem size m . We also experiment with $m = 250$, in this case we have $L \approx 125k$, $|\mathcal{A}| \approx 1.17 \times 10^{149}$, and $R_{\text{Bayes}}(150) \approx 6.56 \times 10^4$, which is only 4.2 times of $R_{\text{Bayes}}(150)$ in the default case. It is worth mentioning that this result also suggests that the Bayes regret bound derived in Theorem 1 is (almost) tight in this problem⁵. To see it, notice that $K = 2m$ and

⁵Recall that Theorem 1 requires $\max_{e \in E} \|\phi_e\|_2 \leq 1$. It can be easily extended to cases with $\max_{e \in E} \|\phi_e\|_2 \leq M$ by scaling the Bayes regret bound by M . However, in this problem ϕ_e is not bounded since it is sampled from a Gaussian distribution. We believe that Theorem 1 can be extended to this case by exploiting

$L = O(m^2)$, and hence the Bayes regret bound derived in Theorem 1 is $\tilde{O}(m)$.

Second, we study how the Bayes cumulative regret of CombLinTS scales with d , the dimension of the feature vectors, by varying $d = 10, 60, 110, \dots, 510$, and demonstrate the result in Figure 1(b). The experiment results indicate that $R_{\text{Bayes}}(150)$ also roughly increases linearly with d , and hence CombLinTS is also scalable with the feature dimension d . This result also suggests that the $\tilde{O}(\sqrt{d})$ bound in Theorem 1 is (almost) tight⁵.

Finally, we study the robustness of CombLinTS with respect to the algorithm parameters σ and λ . In Figure 1(c), we vary $\sigma = 10^{-7}, 10^{-6}, \dots, 10^4$ and in Figure 1(d), we vary $\lambda = 10^{-2}, 10^{-1}, \dots, 10^5$. We would like to emphasize again that we only vary the algorithm parameters and fix $\sigma_{\text{true}} = 1$ and $\lambda_{\text{true}} = 10$. The experiment results show that CombLinTS is robust to the choice of algorithm parameters and performs well for a wide range of σ and λ . However, too small or too large σ , or too small λ , can significantly reduce the performance of CombLinTS, as we have discussed in Section 4.1.

6.2. Online Advertising

In the second experiment, we evaluate CombLinTS on an advertising problem. Our objective is to identify 100 people that are most likely to accept an advertisement offer, subject to the targeting constraint that exactly half of them are females. Specifically, the ground set E includes 33k representative people from Adult dataset (Asuncion & Newman, 2007), which was collected in the 1994 US census. A feasible solution A is any subset of E with $|A| = 100$ and satisfying the targeting constraint mentioned above. We assume that person e accepts an advertisement offer with probability

$$\bar{w}(e) = \begin{cases} 0.15 & \text{income is at least 50k} \\ 0.05 & \text{otherwise,} \end{cases}$$

and people accept offers independently of each other. The features in the generalization matrix Φ are the age, which is binned into 7 groups; gender; whether the person works more than 40 hours per week; and the length of education in years. All these features can be constructed based on the Adult dataset.

CombLinTS is compared to three baselines. The first baseline is the optimal solution A^{opt} . The second baseline is CombUCB1 (Kveton et al., 2015b). This algorithm estimates the probability that person e accepts the offer $\bar{w}(e)$ independently of the other probabilities. The third baseline is CombLinTS without linear generalization, which we sim-

ulate the properties of Gaussian distribution. Roughly speaking, in this problem, with high probability, $\|\phi_e\|_2 = O(\sqrt{d})$.




Figure 2. Regret on the ad dataset.




Figure 3. Regret on the recommendation experiment.

ply refer to as CombTS. As in CombUCB1, this algorithm estimates the probability that person e accepts the offer $\bar{w}(e)$ independently of the other probabilities. The posterior of $\bar{w}(e)$ is modeled as a beta distribution.

Our experiment results are reported in Figure 2. We observe two major trends. First, CombLinTS learns extremely quickly. In particular, its per-step return at episode 100 is 70% of the optimum, and its per-step return at episode 1k is 80% of the optimum. These results are remarkable since the linear generalization is imperfect in this problem. Second, both CombUCB1 and CombTS perform poorly due to insufficient observations with respect to the model complexity. Specifically, in 1k episodes, the people in E are observed 100k times, which implies that each person is observed only 3 times on average. This is not enough to discriminate the people who are likely to accept the advertisement offer from those that are not.

6.3. Artist Recommendation

In the last experiment, we evaluate CombLinTS on a problem of recommending $K = 10$ music artists that are most likely to be chosen by an average user of a music recommendation website. Specifically, the ground set E include artists from the *Last.fm* music recommendation dataset (Cantador et al., 2011). The dataset contains tagging and music artist listening information from a set of 2k users from Last.fm online music system⁶. The tagging part includes the tag assignments of all artists provided by the users. For each user, the artists to whom she listened and the number of listening events are also available in the dataset.

We choose E as the set of artists that were listened by at least two users and had at least one tag assignment among the top 20 most popular tags, and $|E| \approx 6k$. For each artist e , we construct its feature vector $\phi_e \in [0, 1]^{20}$ by setting its

j th component as the fraction of users who assigned tag j to this artist. We assume that each artist e is chosen by an average user with probability $\bar{w}(e) = \frac{1}{|U_e|} \sum_{u \in U_e} w_u(e)$, where U_e is the set of users that listened to artist e , and $w_u(e)$ is the probability that user u likes artist e . We estimate $\bar{w}_u(e)$ based on a Naïve Bayes classifier with respect to the number of person/artist listening events.

Like Section 6.2, we also compare CombLinTS to three baselines: the optimal solution A^{opt} , the CombUCB1 algorithm and the CombTS algorithm. Our experiment results are reported in Figure 3. Similarly as Figure 2, the expected per-step return of CombLinTS approaches that of A^{opt} much faster than CombUCB1 and CombTS. Moreover, both CombUCB1 and CombTS perform poorly due to the insufficient observations with respect to the model complexity: In 1k episodes, each artist is observed less than 2 times on average, which is not enough to discriminate most popular artists from less popular artists.

7. Conclusion

We have proposed two learning algorithms, CombLinTS and CombLinUCB, for stochastic combinatorial semi-bandits with linear generalization. The main contribution of this work is two-fold: First, we have established L -independent regret bounds for these two algorithms under reasonable assumptions, where L is the number of items. Second, we have also evaluated CombLinTS on a variety of problems. The experiment results in the first problem show that CombLinTS is scalable and robust, and the experiment results in the other two problems demonstrate the value of exploiting linear generalization in real-world settings.

It is worth mentioning that our results can be easily extended to the *contextual combinatorial semi-bandits* with linear generalization. In a contextual combinatorial semi-bandit, the probability distribution P (and hence the expected weight \bar{w}) also depends on a context x , which ei-

⁶<http://www.lastfm.com>

ther follows an exogenous stochastic process or is adaptively chosen by an adversary. Assume that each state-item pair (x, e) is associated with a feature vector $\phi_{x,e}$, then similar to Agrawal & Goyal (2013), both CombLinTS and CombLinUCB, as well as their analyses, can be generalized to the contextual combinatorial semi-bandits.

We leave open several questions of interest. One interesting open question is how to derive regret bounds for CombLinTS and CombLinUCB in the agnostic learning cases. Another interesting open question is how to extend the results to combinatorial semi-bandits with nonlinear generalization. We believe that our results can be extended to combinatorial semi-bandits with *generalized linear generalization*⁷, but leave it to future work.

References

- Abbasi-Yadkori, Yasin, Pál, Dávid, and Szepesvári, Csaba. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24*, pp. 2312–2320, 2011.
- Agrawal, Shipra and Goyal, Navin. Analysis of thompson sampling for the multi-armed bandit problem. In *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, pp. 39.1–39.26, 2012.
- Agrawal, Shipra and Goyal, Navin. Thompson sampling for contextual bandits with linear payoffs. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, pp. 127–135, 2013.
- Asuncion, A. and Newman, D.J. UCI machine learning repository, 2007. URL [http://www.ics.uci.edu/\\$\sim\\$mllearn/{MLR}epository.html](http://www.ics.uci.edu/\simmllearn/{MLR}epository.html).
- Audibert, Jean-Yves, Bubeck, Sebastien, and Lugosi, Gábor. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2014.
- Auer, Peter. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3:397–422, 2002.
- Cantador, Iván, Brusilovsky, Peter, and Kuflik, Tsvi. Second workshop on information heterogeneity and fusion in recommender systems (hetrec 2011). In *Proceedings of the 5th ACM conference on Recommender systems, RecSys 2011*. ACM, 2011.
- Cesa-Bianchi, Nicolò and Lugosi, Gábor. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- Chapelle, Olivier and Li, Lihong. An empirical evaluation of Thompson sampling. In *Neural Information Processing Systems*, pp. 2249–2257, 2011.
- Chen, Wei, Wang, Yajun, and Yuan, Yang. Combinatorial multi-armed bandit: General framework and applications. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 151–159, 2013.
- Dani, Varsha, Hayes, Thomas, and Kakade, Sham. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory*, pp. 355–366, 2008.
- Gabillon, Victor, Kveton, Branislav, Wen, Zheng, Eriksson, Brian, and Muthukrishnan, S. Adaptive submodular maximization in bandit setting. In *Advances in Neural Information Processing Systems 26*, pp. 2697–2705, 2013.
- Gabillon, Victor, Kveton, Branislav, Wen, Zheng, Eriksson, Brian, and Muthukrishnan, S. Large-scale optimistic adaptive submodularity. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, 2014.
- Gai, Yi, Krishnamachari, Bhaskar, and Jain, Rahul. Combinatorial network optimization with unknown variables: Multi-armed bandits with linear rewards and individual observations. *IEEE/ACM Transactions on Networking*, 20(5):1466–1478, 2012.
- Koren, Yehuda, Bell, Robert, and Volinsky, Chris. Matrix factorization techniques for recommender systems. *IEEE Computer*, 42(8):30–37, 2009.
- Kveton, Branislav, Wen, Zheng, Ashkan, Azin, and Eydgahi, Hoda. Matroid bandits: Practical large-scale combinatorial bandits. In *Workshops at the Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014a.
- Kveton, Branislav, Wen, Zheng, Ashkan, Azin, Eydgahi, Hoda, and Eriksson, Brian. Matroid bandits: Fast combinatorial optimization with learning. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, pp. 420–429, 2014b.
- Kveton, Branislav, Wen, Zheng, Ashkan, Azin, Eydgahi, Hoda, and Valko, Michal. Learning to act greedily: Polymatroid semi-bandits. *CoRR*, abs/1405.7752, 2014c.
- Kveton, Branislav, Szepesvári, Csaba, Wen, Zheng, and Ashkan, Azin. Cascading bandits: Learning to rank in the cascade model. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015a.

⁷That is, $\bar{w}(e) = h(\phi_e^T \theta^*)$, where $h : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly monotone function.

- Kveton, Branislav, Wen, Zheng, Ashkan, Azin, and Szepesvari, Csaba. Tight regret bounds for stochastic combinatorial semi-bandits. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, 2015b.
- Neu, Gergely and Bartók, Gábor. An efficient algorithm for learning with semi-bandit feedback. In Jain, Sanjay, Munos, Rémi, Stephan, Frank, and Zeugmann, Thomas (eds.), *Algorithmic Learning Theory*, volume 8139 of *Lecture Notes in Computer Science*, pp. 234–248. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-40934-9.
- Papadimitriou, Christos and Steiglitz, Kenneth. *Combinatorial Optimization*. Dover Publications, Mineola, NY, 1998.
- Russo, Daniel and Van Roy, Benjamin. Learning to optimize via posterior sampling. *CoRR*, abs/1301.2609, 2013.
- Russo, Daniel and Van Roy, Benjamin. An information-theoretic analysis of thompson sampling. *CoRR*, abs/1403.5341, 2014.
- Thompson, W.R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Van Roy, Benjamin and Wen, Zheng. Generalization and exploration via randomized value functions. *arXiv preprint arXiv:1402.0635*, 2014.
- Wen, Zheng and Van Roy, Benjamin. Efficient exploration and value function generalization in deterministic systems. In *Advances in Neural Information Processing Systems 26*, pp. 3021–3029, 2013.
- Wen, Zheng, Kveton, Branislav, Eriksson, Brian, and Bhamidipati, Sandilya. Sequential Bayesian search. In *Proceedings of the 30th International Conference on Machine Learning*, pp. 977–983, 2013.
- Yue, Yisong and Guestrin, Carlos. Linear submodular bandits and their application to diversified retrieval. In *Advances in Neural Information Processing Systems 24*, pp. 2483–2491, 2011.

A. Proof for Theorem 1

To prove Theorem 1, we first prove the following theorem:

Theorem 3. *If (1) $\bar{\mathbf{w}} = \Phi\theta^*$, (2) the prior on θ^* is $N(0, \lambda^2 I)$, and (3) the noises are i.i.d. sampled from $N(0, \sigma^2)$, then under CombLinTS algorithm with parameter (Φ, λ, σ) , then we have*

$$R_{\text{Bayes}}(n) \leq 1 + K\lambda \min \left\{ \sqrt{\ln \left(\frac{\lambda L n}{\sqrt{2\pi}} \right)}, \sqrt{d \ln \left(\frac{2dKn\lambda}{\sqrt{2\pi}} \right)} \right\} \sqrt{\frac{2dn \ln \left(1 + \frac{nK\lambda^2}{d} \right)}{\ln \left(1 + \frac{\lambda^2}{\sigma^2} \right)}}. \quad (10)$$

Notice that Theorem 1 follows immediately from Theorem 3. Specifically, if $\lambda \geq \sigma$, then we have

$$\begin{aligned} B_{\text{Bayes}}(n) &\leq 1 + K\lambda \min \left\{ \sqrt{\ln \left(\frac{\lambda L n}{\sqrt{2\pi}} \right)}, \sqrt{d \ln \left(\frac{2dKn\lambda}{\sqrt{2\pi}} \right)} \right\} \sqrt{2dn \log_2 \left(1 + \frac{nK\lambda^2}{d} \right)} \\ &= \tilde{O} \left(K\lambda \sqrt{dn \min \{ \ln(L), d \}} \right). \end{aligned} \quad (11)$$

We now outline the proof of Theorem 3, which is based on (Russo & Van Roy, 2013; Dani et al., 2008). Let \mathcal{H}_t denote the “history” (i.e. all the available information) by the start of episode t . Note that from the Bayesian perspective, conditioning on \mathcal{H}_t , θ^* and θ_t are i.i.d. drawn from $N(\bar{\theta}_t, \Sigma_t)$ (see (Russo & Van Roy, 2013)). This is because that conditioning on \mathcal{H}_t , the posterior belief in θ^* is $N(\bar{\theta}_t, \Sigma_t)$ and based on Algorithm 2, θ_t is independently sampled from $N(\bar{\theta}_t, \Sigma_t)$. Since ORACLE is a fixed combinatorial optimization algorithm (even though it can be independently randomized), and E, \mathcal{A}, Φ are all fixed, then conditioning on \mathcal{H}_t , A^* and A^t are also i.i.d., furthermore, A^* is conditionally independent of θ_t , and A^t is conditionally independent of θ^* .

To simplify the exposition, $\forall \theta \in \mathbb{R}^d$ and $\forall A \subseteq E$, we define

$$g(A, \theta) = \sum_{e \in A} \langle \phi_e, \theta \rangle, \quad (12)$$

then we have $\mathbb{E}[f(A^*, \mathbf{w}_t) | \mathcal{H}_t, \theta^*, \theta_t, A^*, A^t] = g(A^*, \theta^*)$ and $\mathbb{E}[f(A^t, \mathbf{w}_t) | \mathcal{H}_t, \theta^*, \theta_t, A^*, A^t] = g(A^t, \theta^*)$, hence we have $\mathbb{E}[R_t | \mathcal{H}_t] = \mathbb{E}[g(A^*, \theta^*) - g(A^t, \theta^*) | \mathcal{H}_t]$. We also define the *upper confidence bound (UCB)* function $U_t : 2^E \rightarrow \mathbb{R}$ as

$$U_t(A) = \sum_{e \in A} \left[\langle \phi_e, \bar{\theta}_t \rangle + c \sqrt{\phi_e^T \Sigma_t \phi_e} \right], \quad (13)$$

where $c > 0$ is a constant to be specified. Notice that conditioning on \mathcal{H}_t , U_t is a deterministic function and A^*, A^t are i.i.d., then $\mathbb{E}[U_t(A^t) - U_t(A^*) | \mathcal{H}_t] = 0$ and

$$\mathbb{E}[R_t | \mathcal{H}_t] = \mathbb{E}[g(A^*, \theta^*) - U_t(A^*) | \mathcal{H}_t] + \mathbb{E}[U_t(A^t) - g(A^t, \theta^*) | \mathcal{H}_t]. \quad (14)$$

One key observation is that

$$\begin{aligned} \mathbb{E}[U_t(A^t) - g(A^t, \theta^*) | \mathcal{H}_t] &\stackrel{(a)}{=} \sum_{e \in E} \mathbb{E} \left[\mathbb{1}\{e \in A^t\} \left[\langle \phi_e, \bar{\theta}_t - \theta^* \rangle + c \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \middle| \mathcal{H}_t \right] \\ &\stackrel{(b)}{=} \sum_{e \in E} \mathbb{E}[\mathbb{1}\{e \in A^t\} | \mathcal{H}_t] \mathbb{E}[\langle \phi_e, \bar{\theta}_t - \theta^* \rangle | \mathcal{H}_t] + c \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \middle| \mathcal{H}_t \right] \\ &\stackrel{(c)}{=} c \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \middle| \mathcal{H}_t \right], \end{aligned} \quad (15)$$

where (b) follows from the fact that A^t and θ^* are conditionally independent, and (c) follows from $\mathbb{E}[\theta^* | \mathcal{H}_t] = \bar{\theta}_t$. Hence $B_{\text{Bayes}}(n) = \sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] + c \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \right]$. We can show that (1)

$\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] \leq 1$ if we choose

$$c \geq \min \left\{ \sqrt{\ln \left(\frac{\lambda L n}{\sqrt{2\pi}} \right)}, \sqrt{d \ln \left(\frac{2dK n \lambda}{\sqrt{2\pi}} \right)} \right\}, \quad (16)$$

and (2) $\sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \leq K \lambda \sqrt{2dn \ln \left(1 + \frac{nK\lambda^2}{d} \right) / \ln \left(1 + \frac{\lambda^2}{\sigma^2} \right)}$. Thus, the bound in Theorem 3 holds. Please refer to the remainder of this section for the full proof.

A.1. Bound on $\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)]$

We first prove that if we choose

$$c \geq \min \left\{ \sqrt{\ln \left(\frac{\lambda L n}{\sqrt{2\pi}} \right)}, \sqrt{d \ln \left(\frac{2dK n \lambda}{\sqrt{2\pi}} \right)} \right\}, \quad (17)$$

then $\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] \leq 1$. To prove this result, we use the following inequality for truncated Gaussian distribution.

Lemma 1. *If $X \sim N(\mu, s^2)$, then we have*

$$\mathbb{E}[X \mathbb{1}\{X \geq 0\}] = \mu \left[1 - \Phi_G \left(\frac{-\mu}{s} \right) \right] + \frac{s}{\sqrt{2\pi}} \exp \left(-\frac{\mu^2}{2s^2} \right),$$

where Φ_G is the cumulative distribution function (CDF) of the standard Gaussian distribution $N(0, 1)$. Furthermore, if $\mu \leq 0$, we have $\mathbb{E}[X \mathbb{1}\{X \geq 0\}] \leq \frac{s}{\sqrt{2\pi}} \exp \left(-\frac{\mu^2}{2s^2} \right)$.

Based on Lemma 1, we can prove the following lemmas:

Lemma 2. *If $c \geq \sqrt{\ln \left(\frac{\lambda L n}{\sqrt{2\pi}} \right)}$, then we have $\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] \leq 1$.*

Proof. We have the following naive bound:

$$\begin{aligned} g(A^*, \theta^*) - U_t(A^*) &= \sum_{e \in A^*} \left[\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \\ &\leq \sum_{e \in A^*} \left[\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \mathbb{1} \left\{ \langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \geq 0 \right\} \\ &\leq \sum_{e \in E} \left[\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \mathbb{1} \left\{ \langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \geq 0 \right\}. \end{aligned}$$

Notice that conditioning on \mathcal{H}_t , $\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e}$ is a Gaussian random variable with mean $-c \sqrt{\phi_e^T \Sigma_t \phi_e}$ and variance $\phi_e^T \Sigma_t \phi_e$. Thus, from Lemma 1, we have

$$\begin{aligned} &\mathbb{E}_{\theta^*, A^*} [g(A^*, \theta^*) - U_t(A^*) | \mathcal{H}_t] \\ &\stackrel{(a)}{\leq} \sum_{e \in E} \mathbb{E}_{\theta^*} \left[\left[\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \mathbb{1} \left\{ \langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \geq 0 \right\} \middle| \mathcal{H}_t \right] \\ &\stackrel{(b)}{\leq} \sum_{e \in E} \sqrt{\frac{\phi_e^T \Sigma_t \phi_e}{2\pi}} \exp \left(-\frac{c^2}{2} \right) \\ &\stackrel{(c)}{\leq} \exp \left(-\frac{c^2}{2} \right) \sum_{e \in E} \frac{\lambda \|\phi_e\|}{\sqrt{2\pi}} \leq \exp \left(-\frac{c^2}{2} \right) \frac{\lambda L}{\sqrt{2\pi}}, \end{aligned} \quad (18)$$

where the last two inequalities follow from the fact that $\phi_e^T \Sigma_t \phi_e \leq \phi_e^T \Sigma_1 \phi_e \leq \lambda^2 \|\phi_e\|^2 \leq \lambda^2$, since $\|\phi_e\| \leq 1$ by assumption⁸. Thus we have

$$\mathbb{E} \left[\sum_{t=1}^n [g(A^*, \theta^*) - U_t(A^*)] \right] \leq \exp \left(-\frac{c^2}{2} \right) \frac{n\lambda L}{\sqrt{2\pi}}. \quad (19)$$

If we choose $c \geq \sqrt{2 \ln \left(\frac{\lambda L n}{\sqrt{2\pi}} \right)}$, then we have $\mathbb{E}[\sum_{t=1}^n [g(A^*, \theta^*) - U_t(A^*)]] \leq 1$. \square

Lemma 3. *If $c \geq \sqrt{d \ln \left(\frac{2dKn\lambda}{\sqrt{2\pi}} \right)}$, then we also have $\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] \leq 1$.*

Proof. We use v_1, \dots, v_d to denote a fixed set of d orthonormal eigenvectors of Σ_t , and $\Lambda_1^2, \dots, \Lambda_d^2$ to denote the associated eigenvalues. Notice that for $i \neq j$, we have $v_i^T \Sigma_t v_j = \Lambda_i^2 v_i^T v_j = 0$. $\forall i = 1, \dots, d$, we define $v_{i+d} = -v_i$ and $\Lambda_{i+d} = \Lambda_i$, which allows us to define the following conic decomposition:

$$\phi_e = \sum_{i=1}^{2d} \alpha_{ei} v_i, \quad \forall e \in E,$$

subject to the constraints that $\alpha_{ei} \geq 0, \forall (e, i)$. Notice that α_{ei} 's are uniquely determined. Furthermore, for i and j s.t. $|i - j| = d$, by definition of conic decomposition, we have $\alpha_{ei} \alpha_{ej} = 0$. In other words, α_e is a d -sparse vector.

Since we assume that $\|\phi_e\| \leq 1$, we have that $\sum_{i=1}^{2d} \alpha_{ei}^2 \leq 1, \forall e \in E$. Thus, for any e , we have that $\langle \phi_e, \theta^* - \bar{\theta}_t \rangle = \sum_{i=1}^{2d} \alpha_{ei} \langle v_i, \theta^* - \bar{\theta}_t \rangle$ and

$$\begin{aligned} \phi_e^T \Sigma_t \phi_e &= \left(\sum_{i=1}^{2d} \alpha_{ei} v_i^T \right) \Sigma_t \left(\sum_{j=1}^{2d} \alpha_{ej} v_j \right) \\ &= \sum_{i=1}^{2d} \sum_{j=1}^{2d} \alpha_{ei} \alpha_{ej} v_i^T \Sigma_t v_j. \end{aligned} \quad (20)$$

Notice that for $i \neq j$, if $|i - j| \neq d$, then $v_i^T \Sigma_t v_j = 0$; on the other hand, if $|i - j| = d$, $\alpha_{ei} \alpha_{ej} = 0$. Thus, if $i \neq j$, we have $\alpha_{ei} \alpha_{ej} v_i^T \Sigma_t v_j = 0$. Consequently,

$$\phi_e^T \Sigma_t \phi_e = \sum_{i=1}^{2d} \alpha_{ei}^2 v_i^T \Sigma_t v_i = \sum_{i=1}^{2d} \alpha_{ei}^2 \Lambda_i^2.$$

Thus we have

$$\sqrt{\phi_e^T \Sigma_t \phi_e} = \sqrt{\sum_{i=1}^{2d} \alpha_{ei}^2 \Lambda_i^2} \geq \frac{1}{\sqrt{d}} \sum_{i=1}^{2d} \alpha_{ei} \Lambda_i, \quad (21)$$

where the inequality follows from Cauchy-Schwartz inequality, specifically, define $s_i = 1$ if $\alpha_{ei} \Lambda_i \neq 0$, and $s_i = 0$ if $\alpha_{ei} \Lambda_i = 0$, then we have

$$\sum_{i=1}^{2d} \alpha_{ei} \Lambda_i = \sum_{i=1}^{2d} \alpha_{ei} \Lambda_i s_i \leq \sqrt{\sum_{i=1}^{2d} s_i^2} \sqrt{\sum_{i=1}^{2d} \alpha_{ei}^2 \Lambda_i^2} \leq \sqrt{d} \sqrt{\sum_{i=1}^{2d} \alpha_{ei}^2 \Lambda_i^2},$$

⁸Notice that in the derivation of Inequality (18), we implicitly assume that $\phi_e^T \Sigma_t \phi_e > 0, \forall e \in E$. It is worth pointing out that the case with $\phi_e^T \Sigma_t \phi_e = 0$ is a trivial case and this inequality still holds in this case.

where the last inequality follows from the fact that α_e is d -sparse. Thus, for any e , we have that

$$\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c\sqrt{\phi_e^T \Sigma_t \phi_e} \leq \sum_{i=1}^{2d} \alpha_{ei} \langle v_i, \theta^* - \bar{\theta}_t \rangle - \frac{c}{\sqrt{d}} \sum_{i=1}^{2d} \alpha_{ei} \Lambda_i. \quad (22)$$

Consequently, we have

$$\sum_{e \in A^*} \left[\langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c\sqrt{\phi_e^T \Sigma_t \phi_e} \right] \leq \sum_{i=1}^{2d} \left(\langle v_i, \theta^* - \bar{\theta}_t \rangle - \frac{c\Lambda_i}{\sqrt{d}} \right) \left(\sum_{e \in A^*} \alpha_{ei} \right). \quad (23)$$

Define $X_i = \langle v_i, \theta^* - \bar{\theta}_t \rangle - \frac{c\Lambda_i}{\sqrt{d}}$, notice that conditioning on \mathcal{H}_t , we have $X_i | \mathcal{H}_t \sim N\left(-\frac{c\Lambda_i}{\sqrt{d}}, \Lambda_i^2\right)$. Hence we have

$$\begin{aligned} \sum_{e \in A^*} \langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c\sqrt{\phi_e^T \Sigma_t \phi_e} &\stackrel{(a)}{\leq} \sum_{i=1}^{2d} X_i \left[\sum_{e \in A^*} \alpha_{ei} \right] \\ &\stackrel{(b)}{\leq} \sum_{i=1}^{2d} X_i \mathbb{1}\{X_i \geq 0\} \left[\sum_{e \in A^*} \alpha_{ei} \right], \end{aligned}$$

where the inequality (b) follows from the fact that $X_i \leq X_i \mathbb{1}\{X_i \geq 0\}$ and $[\sum_{e \in A^*} \alpha_{ei}] \geq 0$. On the other hand, notice that $|A^*| \leq K$

$$\sum_{e \in A^*} \alpha_{ei} \leq \sqrt{|A^*|} \sqrt{\sum_{e \in A^*} \alpha_{ei}^2} \leq \sqrt{|A^*|} \sqrt{\sum_{e \in A^*} \sum_{j=1}^d \alpha_{ej}^2} \leq \sqrt{|A^*|} \sqrt{\sum_{e \in A^*} 1} = |A^*| \leq K.$$

Since $X_i \mathbb{1}\{X_i \geq 0\} \geq 0$, we have

$$\sum_{e \in A^*} \langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c\sqrt{\phi_e^T \Sigma_t \phi_e} \leq K \sum_{i=1}^{2d} X_i \mathbb{1}\{X_i \geq 0\},$$

notice that the RHS does not include A^* . Hence we have

$$\begin{aligned} \mathbb{E}_{\theta^*} [g(A^*, \theta^*) - U_t(A^*) | \mathcal{H}_t] &= \mathbb{E}_{\theta^*} \left[\sum_{e \in A^*} \langle \phi_e, \theta^* - \bar{\theta}_t \rangle - c\sqrt{\phi_e^T \Sigma_t \phi_e} \middle| \mathcal{H}_t \right] \\ &\leq K \sum_{i=1}^{2d} \mathbb{E}_{\theta^*} [X_i \mathbb{1}\{X_i \geq 0\} | \mathcal{H}_t] \\ &\leq K \sum_{i=1}^{2d} \frac{\Lambda_i}{\sqrt{2\pi}} \exp\left(-\frac{c^2}{2d}\right) \leq \frac{2dK\lambda}{\sqrt{2\pi}} \exp\left(-\frac{c^2}{2d}\right), \end{aligned}$$

where the last inequality follows from the fact that $\Lambda_i \leq \lambda$. Hence we have

$$\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] \leq \frac{2dK\lambda}{\sqrt{2\pi}} \exp\left(-\frac{c^2}{2d}\right),$$

if we choose $c \geq \sqrt{2d \ln\left(\frac{2dK\lambda}{\sqrt{2\pi}}\right)}$, then we have $\sum_{t=1}^n \mathbb{E}[f(A^*, \theta^*) - U_t(A^*)] \leq 1$. \square

Combining the results from Lemma 2 and 3, we have proved that if

$$c \geq \min \left\{ \sqrt{\ln\left(\frac{\lambda Ln}{\sqrt{2\pi}}\right)}, \sqrt{d \ln\left(\frac{2dK\lambda}{\sqrt{2\pi}}\right)} \right\},$$

then $\sum_{t=1}^n \mathbb{E}[g(A^*, \theta^*) - U_t(A^*)] \leq 1$.

A.2. Bound on $\sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \right]$

In this subsection, we derive a bound on $\sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \right]$. Our analysis is motivated by the analysis in (Dani et al., 2008). Specifically, we provide a worst-case bound on $\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e}$, for any realization of random variable \mathbf{w}_t 's, θ_t 's, A^t 's, A^* , and θ^* .

Lemma 4. $\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \leq K \lambda \sqrt{\frac{dn \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right)}{\log \left(1 + \frac{\lambda^2}{\sigma^2} \right)}}$.

Proof. To simplify the exposition, we define

$$z_{t,k} = \sqrt{\phi_{a_k}^T \Sigma_t \phi_{a_k}}. \quad (24)$$

First, notice that Σ_t^{-1} is the Gramian matrix and satisfies

$$\Sigma_{t+1}^{-1} = \Sigma_t^{-1} + \frac{1}{\sigma^2} \sum_{k=1}^{|A^t|} \phi_{a_k} \phi_{a_k}^T. \quad (25)$$

Hence for any t, k , we have that

$$\begin{aligned} \det [\Sigma_{t+1}^{-1}] &\geq \det \left[\Sigma_t^{-1} + \frac{1}{\sigma^2} \phi_{a_k} \phi_{a_k}^T \right] = \det \left[\Sigma_t^{-\frac{1}{2}} \left(I + \frac{1}{\sigma^2} \Sigma_t^{\frac{1}{2}} \phi_{a_k} \phi_{a_k}^T \Sigma_t^{\frac{1}{2}} \right) \Sigma_t^{-\frac{1}{2}} \right] \\ &= \det [\Sigma_t^{-1}] \det \left[I + \frac{1}{\sigma^2} \Sigma_t^{\frac{1}{2}} \phi_{a_k} \phi_{a_k}^T \Sigma_t^{\frac{1}{2}} \right] = \det [\Sigma_t^{-1}] \left(1 + \frac{1}{\sigma^2} \phi_{a_k}^T \Sigma_t \phi_{a_k} \right) \\ &= \det [\Sigma_t^{-1}] \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right). \end{aligned} \quad (26)$$

Hence we have that

$$(\det [\Sigma_{t+1}^{-1}])^{|A^t|} \geq (\det [\Sigma_t^{-1}])^{|A^t|} \prod_{k=1}^{|A^t|} \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right). \quad (27)$$

Remark 1. This is where the extra $O(\sqrt{K})$ factor arises. Notice that this extra factor is purely due to linear generalization. Specifically, if $\Phi = I$, then Σ_t 's and Σ_t^{-1} 's will be diagonal, and we have

$$\det [\Sigma_{t+1}^{-1}] = \det [\Sigma_t^{-1}] \prod_{k=1}^{|A^t|} \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right). \quad (28)$$

Notice that Equation 27 further implies that

$$(\det [\Sigma_{t+1}^{-1}])^K \geq (\det [\Sigma_t^{-1}])^K \prod_{k=1}^{|A^t|} \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right), \quad (29)$$

since $\det [\Sigma_{t+1}^{-1}] \geq \det [\Sigma_t^{-1}]$ and $|A^t| \leq K$. Recall that $\det [\Sigma_1^{-1}] = \left(\frac{1}{\lambda^2} \right)^d$, we have that

$$(\det [\Sigma_{n+1}^{-1}])^K \geq (\det [\Sigma_1^{-1}])^K \prod_{t=1}^n \prod_{k=1}^{|A^t|} \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right) = \frac{1}{\lambda^{2dK}} \prod_{t=1}^n \prod_{k=1}^{|A^t|} \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right). \quad (30)$$

On the other hand, we have

$$\text{trace} [\Sigma_{n+1}^{-1}] = \text{trace} \left[\frac{1}{\lambda^2} I + \frac{1}{\sigma^2} \sum_{t=1}^n \sum_{k=1}^{|A^t|} \phi_{a_k} \phi_{a_k}^T \right] = \frac{d}{\lambda^2} + \frac{1}{\sigma^2} \sum_{t=1}^n \sum_{k=1}^{|A^t|} \|\phi_{a_k}\|^2 \leq \frac{d}{\lambda^2} + \frac{nK}{\sigma^2}, \quad (31)$$

where the last inequality follows from the assumption that $\|\phi_e\| \leq 1, \forall e \in E$ and $|A^t| \leq K$. From the trace-determinant inequality, we have

$$\frac{1}{d} \text{trace} [\Sigma_{n+1}^{-1}] \geq (\det [\Sigma_{n+1}^{-1}])^{\frac{1}{d}},$$

which implies that

$$\left(\frac{1}{\lambda^2} + \frac{nK}{d\sigma^2} \right)^{dK} \geq \left(\frac{1}{d} \text{trace} [\Sigma_{n+1}^{-1}] \right)^{dK} \geq (\det [\Sigma_{n+1}^{-1}])^K \geq \frac{1}{\lambda^{2dK}} \prod_{t=1}^n \prod_{k=1}^{|A^t|} \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right).$$

Taking the logarithm, we have

$$dK \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right) \geq \sum_{t=1}^n \sum_{k=1}^{|A^t|} \log \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right). \quad (32)$$

Notice that $z_{t,k}^2 = \phi_{a_k^t}^T \Sigma_t \phi_{a_k^t}$, hence we have that $0 \leq z_{t,k}^2 \leq \phi_{a_k^t}^T \Sigma_1 \phi_{a_k^t} \leq \lambda^2 \|\phi_{a_k^t}\|^2 \leq \lambda^2$. We have the following technical lemma:

Lemma 5. For any real number $x \in [0, \lambda^2]$, we have $x \leq \frac{\lambda^2}{\log(1 + \frac{\lambda^2}{\sigma^2})} \log(1 + \frac{x}{\sigma^2})$.

Proof. Define $h(x) = \frac{\lambda^2}{\log(1 + \frac{\lambda^2}{\sigma^2})} \log(1 + \frac{x}{\sigma^2}) - x$, thus we only need to prove $h(x) \geq 0$ for $x \in [0, \lambda^2]$. Notice that $h(x)$ is a strictly concave function for $x \geq 0$, and $h(0) = 0, h(\lambda^2) = 0$. From Jensen's inequality, for any $x \in (0, \lambda^2)$, we have $h(x) > 0$. \square

Hence we have that

$$\sum_{t=1}^n \sum_{k=1}^{|A^t|} z_{t,k}^2 \leq \frac{\lambda^2}{\log(1 + \frac{\lambda^2}{\sigma^2})} \sum_{t=1}^n \sum_{k=1}^{|A^t|} \log \left(1 + \frac{z_{t,k}^2}{\sigma^2} \right) \leq \frac{dK\lambda^2 \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right)}{\log(1 + \frac{\lambda^2}{\sigma^2})} \quad (33)$$

Finally, we have that

$$\sum_{t=1}^n \sum_{k=1}^{|A^t|} z_{t,k} \leq \sqrt{nK} \sqrt{\sum_{t=1}^n \sum_{k=1}^{|A^t|} z_{t,k}^2} \leq K\lambda \sqrt{\frac{dn \log(1 + \frac{nK\lambda^2}{d\sigma^2})}{\log(1 + \frac{\lambda^2}{\sigma^2})}}. \quad (34)$$

\square

Recall that the above bound holds for any realization of random variables, thus, we have

$$\mathbb{E} \left[\sum_{t=1}^n [U_t(A^t) - g(A^t, \theta^*)] \right] = c\mathbb{E} \left[\sum_{t=1}^n \sum_{k=1}^{|A^t|} z_{t,k} \right] \leq cK\lambda \sqrt{\frac{dn \log(1 + \frac{nK\lambda^2}{d\sigma^2})}{\log(1 + \frac{\lambda^2}{\sigma^2})}}.$$

With

$$c = \min \left\{ \sqrt{\ln \left(\frac{\lambda Ln}{\sqrt{2\pi}} \right)}, \sqrt{d \ln \left(\frac{2dKn\lambda}{\sqrt{2\pi}} \right)} \right\}, \quad (35)$$

and combining the results in the previous subsection, we have proved Theorem 3.

B. Proof for Theorem 2

We start by writing an alternative formula for Σ_t and $\bar{\theta}_t$. Notice that based on Algorithm 1, we have:

$$\begin{aligned}\Sigma_t^{-1} &= \frac{1}{\lambda^2} I + \frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \phi_{a_k^\tau}^T \\ \Sigma_t^{-1} \bar{\theta}_t &= \frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \mathbf{w}_\tau(a_k^\tau)\end{aligned}\tag{36}$$

Interested readers might refer to Appendix C for the derivation of Equation (36). The proof proceeds as follows. We first construct a confidence set of θ^* based on the “self normalized bound” developed in (Abbasi-Yadkori et al., 2011). Then we derive a regret bound based on Lemma 4 derived above.

B.1. Confidence Set

Our construction of confidence set is motivated by the analysis in (Agrawal & Goyal, 2013). We start by defining some useful notation. Specifically, for any $t = 1, 2, \dots, n$, any $k = 1, 2, \dots, |A^t|$, we define

$$\eta_{t,k} = \mathbf{w}_t(a_k^t) - \bar{\mathbf{w}}(a_k^t).$$

One key observation is that $\eta_{t,k}$ ’s form a Martingale difference sequence (MDS)⁹. Moreover, since $\mathbf{w}_t(a_k^t)$ is bounded in interval $[0, 1]$, $\eta_{t,k}$ ’s are sub-Gaussian with constant $R = 1$. We further define

$$\begin{aligned}V_t &= \frac{\sigma^2}{\lambda^2} I + \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \phi_{a_k^\tau}^T \\ \xi_t &= \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \eta_{\tau,k}\end{aligned}$$

As we will see later, we define V_t and ξ_t to use the “self normalized bound” developed in (Abbasi-Yadkori et al., 2011) (see Theorem 1 of (Abbasi-Yadkori et al., 2011)). Notice that based on the above definition, we have $\Sigma_t^{-1} = \frac{1}{\sigma^2} V_t$, and

$$\bar{\theta}_t - \theta^* = \Sigma_t \left(\frac{1}{\sigma^2} \xi_t - \frac{1}{\lambda^2} \theta^* \right).$$

To see why the second equality holds, notice that

$$\begin{aligned}\Sigma_t^{-1} \bar{\theta}_t &= \frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \left(\phi_{a_k^\tau}^T \theta^* + \eta_{\tau,k} \right) \\ &= \left(\Sigma_t^{-1} - \frac{1}{\lambda^2} I \right) \theta^* + \frac{1}{\sigma^2} \xi_t.\end{aligned}$$

Hence, for any $e \in E$, we have

$$\begin{aligned}|\langle \phi_e, \bar{\theta}_t - \theta^* \rangle| &= \left| \phi_e^T \Sigma_t \left(\frac{1}{\sigma^2} \xi_t - \frac{1}{\lambda^2} \theta^* \right) \right| \\ &\leq \|\phi_e\|_{\Sigma_t} \left\| \frac{1}{\sigma^2} \xi_t - \frac{1}{\lambda^2} \theta^* \right\|_{\Sigma_t} \\ &\leq \|\phi_e\|_{\Sigma_t} \left[\frac{1}{\sigma^2} \|\xi_t\|_{\Sigma_t} + \frac{1}{\lambda^2} \|\theta^*\|_{\Sigma_t} \right],\end{aligned}$$

⁹Note that the notion of “time” is indexed by a pair (t, k) , and follows the lexicographical order.

where the first inequality follows from the Cauchy-Schwarz inequality, and the second inequality follows from the triangular inequality. Notice that

$$\|\theta^*\|_{\Sigma_t} \leq \|\theta^*\|_{\Sigma_1} = \lambda \|\theta^*\|_2,$$

hence we have

$$|\langle \phi_e, \bar{\theta}_t - \theta^* \rangle| \leq \|\phi_e\|_{\Sigma_t} \left[\frac{1}{\sigma^2} \|\xi_t\|_{\Sigma_t} + \frac{1}{\lambda} \|\theta^*\|_2 \right].$$

Moreover, we have

$$\frac{1}{\sigma^2} \|\xi_t\|_{\Sigma_t} = \frac{1}{\sigma^2} \|\xi_t\|_{\sigma^2 V_t^{-1}} = \frac{1}{\sigma} \|\xi_t\|_{V_t^{-1}}.$$

So we have

$$|\langle \phi_e, \bar{\theta}_t - \theta^* \rangle| \leq \|\phi_e\|_{\Sigma_t} \left[\frac{1}{\sigma} \|\xi_t\|_{V_t^{-1}} + \frac{1}{\lambda} \|\theta^*\|_2 \right]. \quad (37)$$

The above inequality always holds. We now provide a high probability bound on $\|\xi_t\|_{V_t^{-1}}$, based on the “self normalized bound” proposed in (Abbasi-Yadkori et al., 2011). From Theorem 1 of (Abbasi-Yadkori et al., 2011), we know for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\|\xi_t\|_{V_t^{-1}} \leq \sqrt{2 \log \left(\frac{\det(V_t)^{1/2} \det(V_1)^{-1/2}}{\delta} \right)} \quad \forall t = 1, 2, \dots$$

Obviously, $\det(V_1) = \left[\frac{\sigma^2}{\lambda^2} \right]^d$, on the other hand, we have

$$[\det(V_t)]^{1/d} \leq \frac{\text{trace}(V_t)}{d} = \frac{\sigma^2}{\lambda^2} + \frac{1}{d} \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \|\phi_{a_k^\tau}\|^2 \leq \frac{\sigma^2}{\lambda^2} + \frac{(t-1)K}{d},$$

where the last inequality follows from the assumption that $\|\phi_e\| \leq 1$. Hence, for $t \leq n$, we have

$$[\det(V_t)]^{1/d} \leq \frac{\sigma^2}{\lambda^2} + \frac{nK}{d}.$$

Thus, with probability at least $1 - \delta$, we have

$$\|\xi_t\|_{V_t^{-1}} \leq \sqrt{d \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right) + 2 \log \left(\frac{1}{\delta} \right)} \quad \forall t = 1, 2, \dots, n.$$

Thus, we have the following lemma:

Lemma 6. For any $\lambda, \sigma > 0$ and any $\delta \in (0, 1)$, with probability at least $1 - \delta$, we have

$$|\langle \phi_e, \bar{\theta}_t - \theta^* \rangle| \leq \|\phi_e\|_{\Sigma_t} \left[\frac{1}{\sigma} \sqrt{d \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right) + 2 \log \left(\frac{1}{\delta} \right)} + \frac{\|\theta^*\|_2}{\lambda} \right], \quad (38)$$

for all $t = 1, 2, \dots, n$, and for all $e \in E$.

Notice that $\|\phi_e\|_{\Sigma_t} = \sqrt{\phi_e^T \Sigma_t \phi_e}$, thus, the above lemma immediately implies the following lemma:

Lemma 7. For any $\lambda, \sigma > 0$, any $\delta \in (0, 1)$, and any

$$c \geq \frac{1}{\sigma} \sqrt{d \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right) + 2 \log \left(\frac{1}{\delta} \right)} + \frac{\|\theta^*\|_2}{\lambda},$$

with probability at least $1 - \delta$, we have

$$\langle \phi_e, \bar{\theta}_t \rangle - c \sqrt{\phi_e^T \Sigma_t \phi_e} \leq \langle \phi_e, \theta^* \rangle \leq \langle \phi_e, \bar{\theta}_t \rangle + c \sqrt{\phi_e^T \Sigma_t \phi_e},$$

for all $e \in E$ and $t = 1, 2, \dots, n$.

Notice that

$$\langle \phi_e, \theta^* \rangle \leq \langle \phi_e, \bar{\theta}_t \rangle + c\sqrt{\phi_e^T \Sigma_t \phi_e}$$

is exactly $\bar{\mathbf{w}}(e) \leq \hat{\mathbf{w}}_t(e)$.

B.2. Regret Analysis

We define event G as

$$G = \left\{ \langle \phi_e, \bar{\theta}_t \rangle - c\sqrt{\phi_e^T \Sigma_t \phi_e} \leq \langle \phi_e, \theta^* \rangle \leq \langle \phi_e, \bar{\theta}_t \rangle + c\sqrt{\phi_e^T \Sigma_t \phi_e} \forall e \in E, \forall t = 1, \dots, n \right\}, \quad (39)$$

and use \bar{G} to denote the complement of event G . Recall that Lemma 7 states that if

$$c \geq \frac{1}{\sigma} \sqrt{d \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right) + 2 \log \left(\frac{1}{\delta} \right) + \frac{1}{\lambda} \|\theta^*\|_2}, \quad (40)$$

then $\mathbb{P}(G) \geq 1 - \delta$. Moreover, by definition, under event G , we have $\bar{\mathbf{w}}(e) \leq \hat{\mathbf{w}}_t(e)$, for all $t = 1, \dots, n$ and any $e \in E$.

Notice that

$$\begin{aligned} R(n) &= \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \mathbf{w}_t(e) - \sum_{e \in A^t} \mathbf{w}_t(e) \right] \\ &= \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \right] \\ &= \mathbb{P}(G) \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] + \mathbb{P}(\bar{G}) \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| \bar{G} \right] \\ &\leq \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] + \mathbb{P}(\bar{G}) nK, \end{aligned}$$

where the last inequality follows from the naive bound on the realized regret. If c satisfies inequality (40), we have $\mathbb{P}(\bar{G}) \leq \delta$, hence we have

$$R(n) \leq \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] + nK\delta.$$

Finally, we bound $\sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right]$ using a worst-case bound conditioning on G (worst-case over all the possible random realizations), notice that conditioning on G , we have

$$\sum_{e \in A^*} \bar{\mathbf{w}}(e) \leq \sum_{e \in A^*} \hat{\mathbf{w}}_t(e) \leq \sum_{e \in A^t} \hat{\mathbf{w}}_t(e),$$

where the first inequality follows from the definition of event G , and the second inequality follows from that A^t is the *exact* solution of the combinatorial optimization problem $(E, \mathcal{A}, \hat{\mathbf{w}}_t)$. Thus we have

$$\begin{aligned} \sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) &\leq \sum_{e \in A^t} \hat{\mathbf{w}}_t(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \\ &= \sum_{e \in A^t} \left[\langle \phi_e, \bar{\theta}_t - \theta^* \rangle + c\sqrt{\phi_e^T \Sigma_t \phi_e} \right] \\ &\leq 2c \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e}, \end{aligned}$$

where the last inequality follows from the definition of G . Recall that from Lemma 4, we have

$$\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \leq K\lambda \sqrt{\frac{dn \log \left(1 + \frac{nK\lambda^2}{d\sigma^2}\right)}{\log \left(1 + \frac{\lambda^2}{\sigma^2}\right)}}.$$

Thus we have

$$\sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^*} \bar{\mathbf{w}}(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] \leq 2c\mathbb{E} \left[\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \leq 2cK\lambda \sqrt{\frac{dn \log \left(1 + \frac{nK\lambda^2}{d\sigma^2}\right)}{\log \left(1 + \frac{\lambda^2}{\sigma^2}\right)}},$$

which implies

$$R(n) \leq 2cK\lambda \sqrt{\frac{dn \log \left(1 + \frac{nK\lambda^2}{d\sigma^2}\right)}{\log \left(1 + \frac{\lambda^2}{\sigma^2}\right)}} + nK\delta.$$

C. Technical Lemma

In this section, we derive Equation (36). We first prove the following technical lemma:

Lemma 8. For any $\phi, \bar{\theta} \in \mathbb{R}^d$, any positive definite $\Sigma \in \mathbb{R}^{d \times d}$, any $\sigma > 0$, and any $w \in \mathbb{R}$, if we define

$$\begin{aligned} \Sigma_{\text{new}} &= \Sigma - \frac{\Sigma \phi \phi^T \Sigma}{\phi^T \Sigma \phi + \sigma^2} \\ \bar{\theta}_{\text{new}} &= \left[I - \frac{\Sigma \phi \phi^T}{\phi^T \Sigma \phi + \sigma^2} \right] \bar{\theta} + \left[\frac{\Sigma \phi}{\phi^T \Sigma \phi + \sigma^2} \right] w, \end{aligned}$$

then we have

$$\Sigma_{\text{new}}^{-1} = \Sigma^{-1} + \frac{1}{\sigma^2} \phi \phi^T \tag{41}$$

$$\Sigma_{\text{new}}^{-1} \bar{\theta}_{\text{new}} = \Sigma^{-1} \bar{\theta} + \frac{1}{\sigma^2} \phi w. \tag{42}$$

Proof. Notice that Equation (41) follows directly from the Woodbury matrix identity (matrix inversion lemma). We now prove Equation (42). Notice that we have

$$\begin{aligned} \bar{\theta}_{\text{new}} &= \left[I - \frac{\Sigma \phi \phi^T}{\phi^T \Sigma \phi + \sigma^2} \right] \bar{\theta} + \left[\frac{\Sigma \phi}{\phi^T \Sigma \phi + \sigma^2} \right] w \\ &= \left[\Sigma - \frac{\Sigma \phi \phi^T \Sigma}{\phi^T \Sigma \phi + \sigma^2} \right] \Sigma^{-1} \bar{\theta} + \left[\frac{\Sigma \phi}{\phi^T \Sigma \phi + \sigma^2} \right] w \\ &= \Sigma_{\text{new}} \Sigma^{-1} \bar{\theta} + \left[\frac{\Sigma \phi}{\phi^T \Sigma \phi + \sigma^2} \right] w, \end{aligned}$$

that is,

$$\Sigma_{\text{new}}^{-1} \bar{\theta}_{\text{new}} = \Sigma^{-1} \bar{\theta} + \left[\frac{\Sigma_{\text{new}}^{-1} \Sigma \phi}{\phi^T \Sigma \phi + \sigma^2} \right] w. \tag{43}$$

Notice that

$$\Sigma_{\text{new}}^{-1} \Sigma \phi = \left[\Sigma^{-1} + \frac{1}{\sigma^2} \phi \phi^T \right] \Sigma \phi = \phi + \frac{\phi^T \Sigma \phi}{\sigma^2} \phi = \frac{\sigma^2 + \phi^T \Sigma \phi}{\sigma^2} \phi. \tag{44}$$

Plug Equation (44) into Equation (43), we have Equation (42). \square

Based on Lemma 8, by mathematical induction, we have

$$\begin{aligned}\Sigma_t^{-1} &= \Sigma_1^{-1} + \frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \phi_{a_k^\tau}^T \\ \Sigma_t^{-1} \bar{\theta}_t &= \Sigma_1^{-1} \bar{\theta}_1 + \frac{1}{\sigma^2} \sum_{\tau=1}^{t-1} \sum_{k=1}^{|A^\tau|} \phi_{a_k^\tau} \mathbf{w}_\tau(a_k^\tau),\end{aligned}$$

further noting that $\Sigma_1 = \lambda^2 I$ and $\bar{\theta}_1 = 0$, we can derive Equation (36).

D. A Variant of Theorem 2 for Approximation Algorithms

By suitably redefining the realized regret, we can prove a variant of Theorem 2 in which ORACLE can be an approximation algorithm. Specifically, for a (possibly approximation) algorithm ORACLE, let $A^*(\mathbf{w})$ be the solution of ORACLE to the optimization problem $(E, \mathcal{A}, \mathbf{w})$, we say $\gamma \in [0, 1)$ is a *sub-optimality gap* of ORACLE if

$$f(A^*(\mathbf{w}), \mathbf{w}) \geq (1 - \gamma) \max_{A \in \mathcal{A}} f(A, \mathbf{w}), \quad \forall \mathbf{w}. \quad (45)$$

Then we define the (scaled) realized regret R_t^γ as

$$R_t^\gamma = f(A^{\text{opt}}, \mathbf{w}_t) - \frac{f(A^t, \mathbf{w}_t)}{1 - \gamma}, \quad (46)$$

where A^{opt} is the exact solution to the optimization problem $(E, \mathcal{A}, \bar{\mathbf{w}})$. The (scaled) cumulative regret $R^\gamma(n)$ is defined as

$$R^\gamma(n) = \sum_{t=1}^n \mathbb{E}[R_t^\gamma | \bar{\mathbf{w}}].$$

Under the assumptions that (1) the support of P is a subset of $[0, 1]^L$ (i.e. $\mathbf{w}_t(e) \in [0, 1] \forall t$ and $\forall e \in E$), and (2) the oracle ORACLE has sub-optimality gap $\gamma \in [0, 1)$, we have the following variant of Theorem 2 when CombLinUCB is applied to coherent learning cases:

Theorem 4. *For any $\lambda, \sigma > 0$, any $\delta \in (0, 1)$, and any c satisfying*

$$c \geq \frac{1}{\sigma} \sqrt{d \ln \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right) + 2 \ln \left(\frac{1}{\delta} \right) + \frac{\|\theta^*\|_2}{\lambda}}, \quad (47)$$

if $\bar{\mathbf{w}} = \Phi \theta^$ and the above two assumptions hold, then under CombLinUCB algorithm with parameter $(\Phi, \lambda, \sigma, c)$, we have*

$$R^\gamma(n) \leq \frac{2cK\lambda}{1 - \gamma} \sqrt{\frac{dn \ln \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right)}{\ln \left(1 + \frac{\lambda^2}{\sigma^2} \right)}} + nK\delta.$$

Proof. Notice that Lemma 7 in Section B.1 still holds. With G defined in Equation (39), we have

$$\begin{aligned}R^\gamma(n) &= \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \mathbf{w}_t(e) - \frac{1}{1 - \gamma} \sum_{e \in A^t} \mathbf{w}_t(e) \right] \\ &= \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1 - \gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) \right] \\ &= \mathbb{P}(G) \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1 - \gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] + \mathbb{P}(\bar{G}) \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1 - \gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| \bar{G} \right] \\ &\leq \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1 - \gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] + \mathbb{P}(\bar{G}) nK,\end{aligned}$$

where the last inequality follows from the naive bound on R_t^γ . If c satisfies inequality (40), we have $\mathbb{P}(\bar{G}) \leq \delta$, hence we have

$$R^\gamma(n) \leq \sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1-\gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right] + nK\delta.$$

Finally, we bound $\sum_{t=1}^n \mathbb{E} \left[\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1-\gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) \middle| G \right]$ using a worst-case bound conditioning on G (worst-case over all the possible random realizations), notice that conditioning on G , we have

$$\sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) \leq \sum_{e \in A^{\text{opt}}} \hat{\mathbf{w}}_t(e) \leq \max_{A \in \mathcal{A}} \sum_{e \in A} \hat{\mathbf{w}}_t(e) \leq \frac{1}{1-\gamma} \sum_{e \in A^t} \hat{\mathbf{w}}_t(e),$$

where

- The first inequality follows from the definition of event G . Specifically, under event G , $\bar{\mathbf{w}}(e) \leq \hat{\mathbf{w}}_t(e)$ for all $t = 1, \dots, n$ and all $e \in E$.
- The second inequality follows from $A^{\text{opt}} \in \mathcal{A}$.
- The last inequality follows from $A^t \leftarrow \text{ORACLE}(E, \mathcal{A}, \hat{\mathbf{w}}_t)$ and ORACLE has sub-optimality gap γ (see Equation (45)).

Thus we have

$$\begin{aligned} \sum_{e \in A^{\text{opt}}} \bar{\mathbf{w}}(e) - \frac{1}{1-\gamma} \sum_{e \in A^t} \bar{\mathbf{w}}(e) &\leq \frac{1}{1-\gamma} \left[\sum_{e \in A^t} \hat{\mathbf{w}}_t(e) - \sum_{e \in A^t} \bar{\mathbf{w}}(e) \right] \\ &= \frac{1}{1-\gamma} \sum_{e \in A^t} \left[\langle \phi_e, \bar{\theta}_t - \theta^* \rangle + c \sqrt{\phi_e^T \Sigma_t \phi_e} \right] \\ &\leq \frac{2c}{1-\gamma} \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e}, \end{aligned}$$

where the last inequality follows from the definition of G . Recall that from Lemma 4, we also have

$$\sum_{t=1}^n \sum_{e \in A^t} \sqrt{\phi_e^T \Sigma_t \phi_e} \leq K\lambda \sqrt{\frac{dn \log \left(1 + \frac{nK\lambda^2}{d\sigma^2} \right)}{\log \left(1 + \frac{\lambda^2}{\sigma^2} \right)}}.$$

Putting the above inequalities together, we have proved the theorem. \square